

VARIATIONAL DATA ASSIMILATION FOR TRANSIENT BLOOD FLOW SIMULATIONS

S. W. FUNKE*, M. NORDAAS†, Ø. EVJU‡, M. S. ALNÆS§, AND K.-A. MARDAL¶

Abstract. Several cardiovascular diseases are caused from localised abnormal blood flow such as in the case of stenosis or aneurysms. Prevailing theories propose that the development is caused by abnormal wall-shear stress in focused areas. Computational fluid mechanics have arisen as a promising tool for a more precise and quantitative analysis, in particular because the anatomy is often readily available even by standard imaging techniques such as magnetic resonance and computed tomography angiography. However, computational fluid mechanics rely on accurate boundary conditions which is difficult to obtain. In this paper we address the problem of recovering high resolution information from noisy, low-resolution measurements of blood flow using variational data assimilation (also known as 4DVar). We show that accurate flow reconstruction is obtained with proper regularisation even in the presence of significant noise and for a range of regularisation parameters spanning orders of magnitude. Numerical experiments are performed in both 2D and 3D and with pulsatile flow relevant for physiological flow in cerebral aneurysms.

Key words. blood flow, variational data assimilation, finite element method, adjoint equations, Navier-Stokes, BFGS

AMS subject classifications. 35Q92, 35Q93, 65K10, 76D55, 35Q30

1. Introduction. Detailed insight of blood flow has the potential to assist clinical decisions, for example when evaluating the risk of rupture of an aneurysm [5, 31, 28]. Different non-invasive measurement techniques for blood flow exist today, such as ultrasound or phase-contrast magnetic resonance imaging. Unfortunately, the spatial and temporal resolution of these techniques are still too coarse to unveil potentially important flow details. Computational patient-specific blood flow models are promising tools for obtaining blood flow information with nearly arbitrary high temporal and spatial resolution. In addition, they allow for computation of non-observable variables, such as the blood pressure and wall-shear stresses, which are considered important factors in vascular diseases [28, 25, 7, 18, 25]. The validity of such simulations depends on the accuracy of the inflow and outflow velocities at the open vessel boundaries and the segmentation of the vascular geometry. These model parameters are typically patient-specific and partially or fully unknown [24, 23].

It then seems obvious to incorporate physical measurements into the model in order to identify the unknown parameters. The result would be a high-resolution blood flow simulation that best matches the available measurements. This idea, known as variational data assimilation, has been successfully applied to weather and ocean modelling, see for example [11], and to a more general setting for the optimal control of the Navier-Stokes equations [10, 19, 12]. In the context of blood-flow simulations, this technique has been applied to the steady-state Navier-Stokes equations, for example in [29, 19, 13, 12, 2, 1, 6, 16]. However, the steady-state assumption is inadequate for the pulsatile blood flow in larger arteries.

In this paper we consider the variational data assimilation for *transient* blood flow models. Section 2 formulates the data-assimilation problem as a mathematical

*Center for Biomedical Computing, Simula Research Laboratory, Norway (simon@simula.no)

†Center for Biomedical Computing, Simula Research Laboratory, Norway

‡Center for Biomedical Computing, Simula Research Laboratory, Norway

§Center for Biomedical Computing, Simula Research Laboratory, Norway

¶Department of Mathematics, University of Oslo, Norway




FIG. 1. The model scenario considered in this paper: a small subset of the artery system with one inlet and two outlet boundaries in 2D and 3D.

optimisation problem constrained by the Navier-Stokes equations. Section 3 discusses the numerical details. For the optimisation method, focus will be on achieving convergence which is independent of the choice of discretisation for the Navier-Stokes equations. Furthermore, special considerations will be put on the inclusion of data that are coarse with respect to the time resolution because the number of samples per cardiac cycle is typically in the order 20-40 while the number of time steps in a CFD simulation typically is 100-10000. Section 4 demonstrate the feasibility of this approach through numerical examples in two and three dimensions. This is the first study where variational data assimilation methods have been used in a three-dimensional transient blood flow solver.

2. Mathematical formulation.

2.1. Blood flow model. Most computational modelling in cerebral aneurysm studies assume Newtonian flow with rigid walls, which appear to be adequate [27, 8]. Therefore, we model the blood flow through a vessel with the incompressible Navier-Stokes equations

$$(1) \quad \begin{aligned} u_t + (u \cdot \nabla)u - \nu \Delta u + \nabla p &= f & \text{in } \Omega \times (0, T], \\ \nabla \cdot u &= 0 & \text{in } \Omega \times (0, T]. \end{aligned}$$

Here, $\Omega \times (0, T]$ is the space-time domain, u and p are the blood velocity and (scaled) pressure fields, ν is the (kinematic) viscosity and f describes external body forces. A more complete blood flow model could incorporate non-Newtonian effects and the fluid structure interactions between the blood and the vessel wall [30]. However, for the purpose of this paper it is sufficient to consider (1) and to note that the proposed techniques also apply to more complex models.

We only model a small subset of the artery system and the boundaries of the computational domain consists of a physical boundary, the vessel walls, as well nonphysical boundaries at inlets and outlets. Again for simplicity, we consider the common scenario of one inlet and two outlets as sketched in figure 1. To close the system, we specify suitable initial and boundary conditions

$$\begin{aligned} (2a) \quad & u = u_0, & \text{on } \Omega \times \{t = 0\}, \\ (2b) \quad & u = g_D, & \text{on } \Gamma_D \times (0, T], \\ (2c) \quad & pn - \mu \partial_n u = 0, & \text{on } \Gamma_{\text{out}_2} \times (0, T], \\ (2d) \quad & u = 0, & \text{on } \Gamma_{\text{walls}} \times (0, T]. \end{aligned}$$

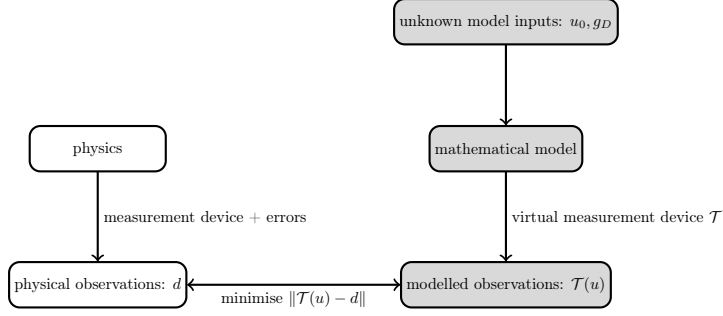


FIG. 2. Variational data assimilation replicates the observation steps in a mathematical model and minimises the discrepancy between measured and modelled observations by varying the model inputs.

with normal vector n and a Dirichlet boundary $\Gamma_D := \Gamma_{\text{in}} \cup \Gamma_{\text{out}_1}$. A traction free boundary is assumed on the outlet Γ_{out_2} , which implies that the vessel is straight in the surroundings of this outlet.

2.2. Variational data assimilation. Variational data assimilation is a technique to recover unknown data from given observation. The idea is to build a model that replicates the steps of the measurement acquisition, and to tweak the free model parameter so that the discrepancy between observed and modelled measurements is minimised (figure 2). Variational data assimilation for time-dependent problems, or 4DVar, has been successfully applied to problems in meteorology and oceanography and is used in production services for weather and ocean forecasts and retrospective analysis.

In the blood flow problem, the aim of variational data assimilation is to recover flow velocity and pressure fields from observational data that is typically noisy, as well as limited in spatial and temporal resolution. These recovered fields should be physically reasonable, in the sense that they satisfy the mathematical model considered. More specifically, the initial condition u_0 and Dirichlet boundary condition g_D in (2) have to be determined in order to solve the model equations (1). In the data assimilation setting these fields are unknown. Instead, one has some given observational data d , and aim to recover the fields u_0 and g_D that best reproduces the data d . In the present paper, it is assumed that the data is a set of N measured velocity fields on a subdomain $\Omega_{\text{obs}} \subset \Omega$. To be precise, it is assumed that $d = (d_1, \dots, d_N) \in L^2(\Omega_{\text{obs}})^N$.

The objective is to recover the initial and boundary conditions by minimising the misfit between simulation and measurement data. Hence we define a goal quantity

$$(3) \quad J(u) = \|\mathcal{T}u - d\|^2 = \sum_{n=1}^N \int_{\Omega_{\text{obs}}} |\mathcal{T}_n u - d_n|^2 dx,$$

where \mathcal{T}_n are the observation operators which approximate the physical measurement device by mapping the velocity solution u to simulated measurements. For instance, if the measurement device takes instantaneously measurements at N timelevels t_1, \dots, t_N , then the observation operators should be pointwise evaluations in time. Other common choices for the observation operators are time and space averaged evaluations of the velocity state. The observations might be available only in parts of the domain, hence the functional integrates over $\Omega_{\text{obs}} \subseteq \Omega$.

We can now formulate the data assimilation problem as an optimisation problem constrained by the Navier-Stokes equations:

$$(4) \quad \min_{\substack{(u,p) \in Y \\ (u_0, g_D) \in M}} J(u) + \mathcal{R}(u_0, g_D) \quad \text{subject to (1)-(2),}$$

where M and Y are suitable function spaces, to be determined later in section 3.3 below. The Tikhonov regularisation \mathcal{R} enforces smoothness of the controls:

$$(5) \quad \mathcal{R}(u_0, g_D) = \frac{\alpha}{2} \|g_D\|_{\Gamma_D \times (0, T]}^2 + \frac{\gamma}{2} \|u_0\|_{\Omega}^2,$$

where the coefficients α and γ determine how strongly the problem is regularised in the given norms.

2.2.1. Choice of norms. The choice of norms in the regularisation term (5) specifies the expected regularity of the reconstructed blood flow. For instance, [16] has shown that a unsuitable choice can have a negative impact on the quality of the reconstructed data.

The norm used for the initial data is

$$\|u_0\|_{\Omega} = \|u_0\|_{H^1(\Omega)} = \left(\int_{\Omega} |u_0|^2 + |\nabla u_0|^2 dx \right)^{\frac{1}{2}},$$

and for the boundary

$$\|g_D\|_{\Gamma_D \times (0, T]} = \left(\int_0^T \int_{\Omega} |g_D|^2 + |\nabla g_D|^2 + |\dot{g}_D|^2 + |\nabla \dot{g}_D|^2 dx dt \right)^{\frac{1}{2}}.$$

We remark that the norms used require more smoothness on the boundary and initial data than is usually required for the variational formulation of the Navier-Stokes equations, in particular for the time derivative \dot{g}_D .

3. Numerical solution.

3.1. Formulation of the reduced problem. Multiple strategies exist for solving the data assimilation problems. One option is to derive and solve the first order optimality system of (4). This leads to a large, non-linear system that couples all spatial and temporal degrees of freedoms of the discretised Navier-Stokes and adjoint Navier-Stokes equations. Solving this system is numerically challenging and requires the development of specialised solvers.

The approach taken here is based on the reduced optimisation problem of (4). The reduced problem is formed by considering the velocity solution as an implicit function of the initial and boundary controls by solving the Navier-Stokes equations (1) and (2). We denote this velocity operator as $u(u_0, g_D)$.

To simplify notation, let $m = (u_0, g_D) \in M$ denote the controlled variable. The functional (3) now has the reduced form

$$(6) \quad \hat{J}(m) := J(u(m)) + \mathcal{R}(m).$$

The *reduced optimisation problem* reads

$$(7) \quad \min_{m \in M} \hat{J}(m).$$

Note that in contrast to (4), the reduced problem is an unconstrained optimisation problem. As a consequence, it can be solved with established unconstrained optimisation methods. Also, evaluating the reduced functional requires the solution of a Navier-Stokes system. Here, standard solver techniques can be directly applied.

3.2. Optimisation . The reduced minimisation problem (7) is solved with the Broyden-Fletcher-Goldbarb-Shanno (BFGS) algorithm. In this section we present brief overview of the method and its implementation.

The minimisation problem (7) is iteratively solved by generating a sequence of points m_0, m_1, \dots , approximating a miniser of \hat{J} . In each iteration, evaluations of the derivative $D\hat{J}(m_k) \in M^*$ are used to determine a direction $d_k \in M$ in which the functional is decreasing. This general descent algorithm in Hilbert spaces is formulated in algorithm 1.

Algorithm 1 A general descent algorithm in Hilbert spaces, applied to the reduced minimisation problem (7).

```

Choose an initial point  $x_0$ 
for  $x = 0, 1, \dots$  do
    Choose a search direction  $d_k = -H_k D\hat{J}(x_k)$ 
    Choose a step length  $\alpha_k > 0$  such that  $\hat{J}(x_k + \alpha_k d_k) < \hat{J}(x_k)$ 
    Set  $x_{k+1} = x_k + \alpha_k d_k$ 
    if converged then
        return
    end if
end for

```

The search direction $d_k = H_k D\hat{J}(m_k)$ is a descent direction if the operator $H_k : M^* \rightarrow M$ is positive definite and self-adjoint. The choice of operators H_k mapping derivatives to search directions essentially characterises the method. For example, taking $H_k = H$ as the Riesz operator for M (i.e. choosing d_k to be the gradient of \hat{J} at m_k) results in a steepest algorithm. Setting $H_k = D^2\hat{J}(m_k)^{-1}$, assuming \hat{J} is convex, results in a Newton algorithm.

In the present paper, the algorithm used is the Broyden-Fletcher-Goldbarb-Shanno (BFGS) algorithm, which is a descent method of quasi-Newton type. Quasi-Newton methods have good convergence properties and do not require evaluations of the Hessian. Instead, such methods maintains an iteratively constructed approximation to the inverse of the Hessian. The update formula specific to the BFGS algorithm is

$$(8) \quad H_{k+1} = \left(1 - \frac{s_{k+1} \otimes y_{k+1}}{\rho_{k+1}}\right) H_k \left(1 - \frac{y_{k+1} \otimes s_{k+1}}{\rho_{k+1}}\right) + \frac{s_{k+1} \otimes s_{k+1}}{\rho_{k+1}}$$

see e.g. [22, Chapter 6]. Here, $\otimes : X \times Y \rightarrow \mathcal{B}(Y^*, X)$ denotes the outer product defined by $(x \otimes y)(z) = x\langle z, y \rangle_{Y^*, Y}$, for $x \in X$ and $y \in Y$, where $\langle \cdot, \cdot \rangle_{Y^*, Y}$ denotes duality coupling, and

$$\begin{aligned} s_{k+1} &= m_{k+1} - m_k, \\ y_{k+1} &= D\hat{J}(m_{k+1}) - D\hat{J}(m_k), \\ \rho_{k+1} &= \langle y_{k+1}, s_k \rangle_{M^*, M}. \end{aligned}$$

Note that the initial $H_0 : M^* \rightarrow M$ has to be prescribed. A natural choice is to take H_0 to be Riesz operator for the space M . That is, H_0 is the unique operators such that

$$(9) \quad (m_0, m_1)_M = \langle H_0^{-1} m_0, m_1 \rangle_{M^*, M}$$

for all $m_0, m_1 \in M$. This definition of H_0 allows for mesh-independent convergence [26], and is readily seen to coincide with the second order partial derivative of \hat{J} with

respect to M . If $D^2\hat{J}(m) - H_0^{-1}$ is compact, the method converges superlinearly, see e.g. [17].

For practical implementations, it is common to truncate the update formula (8) and store only the last 3 – 10 pairs of vectors y_k and s_k . The step lengths α_k in algorithm 1 should be chosen to satisfy the Wolfe conditions, which ensures the convergence of the method [22, chapter 6].

3.3. Discretisation. The optimisation method in section 3.2 requires evaluations of the reduced functional $\hat{J}(m)$ and its derivative $D\hat{J}(m)$. Evaluating the reduced functional requires the numerical solution of the Navier-Stokes equations. This is described in section 3.3.1. The derivatives are computed by solving the adjoint equations, described in section 3.3.2.

3.3.1. Discretisation of the Navier-Stokes equations. The Navier-Stokes equations are discretised with a θ time-stepping scheme and the finite element method. The controlled Dirichlet boundary conditions are weakly enforced with a variant of the Nitsche method [21]. An advantage of the Nitsche approach is that the boundary values are explicitly included in the variational formulation, which simplifies the (automated) derivation of the adjoint equations. This is exploited in the implementation.

For the spatial discretisation, we consider conforming finite element spaces

$$(10) \quad \begin{aligned} V_h &\subset H_{0,\Gamma_{\text{walls}}}^1(\Omega) = \{u \in H^1(\Omega) : u|_{\Gamma_{\text{walls}}} = 0\} \\ Q_h &\subset L^2(\Omega). \end{aligned}$$

For the time discretisation we assume a partition of the interval $[0, T]$ with a constant timestep δt . Applying a standard θ time-stepping scheme to the Navier-Stokes equations (1), we obtain a sequence of nonlinear problems: For $k = 0, \dots, N-1$, let $u^{k+\theta} = \theta u^{k+1} + (1-\theta)u^k$ and find $(u^{k+1}, p^{k+1}) \in V_h \times Q_h$ such that

$$(11) \quad \begin{aligned} \frac{u^{k+1} - u^k}{\delta t} - \nu \Delta u^{k+\theta} + (u^{k+\theta} \cdot \nabla) u^{k+\theta} - \nabla p^{k+1} &= 0, \\ \nabla \cdot u^{k+1} &= 0, \end{aligned}$$

subject to the boundary conditions (2). The equations (11) are integrated against test functions $v \in V_h$ and $q \in Q_h$ in order to obtain a nonlinear variational problem at each time t_k ,

$$(12) \quad \begin{aligned} 0 &= \int_{\Omega} \left(\frac{u^{k+1} - u^k}{\delta t} \right) \cdot v + \nu \nabla u^{k+\theta} : \nabla v \, dx \\ &\quad + \int_{\Omega} (u^{k+\theta} \cdot \nabla) u^{k+\theta} \cdot v \, dx \\ &\quad + \int_{\Omega} q \nabla \cdot u^{k+1} + p^{k+1} \nabla \cdot v \, dx \\ &\quad - \int_{\Gamma_D} \left(\nu \frac{\partial u^{k+\theta}}{\partial n} - p^{k+1} n \right) \cdot v \, ds \\ &\quad - \int_{\Gamma_D} \left(\theta \nu \frac{\partial v}{\partial n} - q n \right) \cdot (u^{k+1} - g^{k+1}) \, ds \\ &\quad + \int_{\Gamma_D} \frac{\nu \sigma}{h} (u^{k+1} - g^{k+1}) \cdot v \, ds. \end{aligned}$$

The nonlinear variational problem (12) consists of a volume integral and a boundary integral over Γ_D . The volume integral coincides with the “standard” variational form of (11) obtained when the boundary condition (2b) is strongly imposed. The second, boundary integral part of the variational problem arises from the weakly imposing the Dirichlet boundary condition (15) with Nitsche’s method, and is discussed in detail below.

The discrete spaces for the state and control variables are

$$\begin{aligned} Y &= V_h^N \times Q_h^N \\ M &= V_h \times (T_\Gamma V_h)^N, \end{aligned}$$

and we introduce the notation

$$\begin{aligned} y &= (u, p) \in Y, \\ u &= (u_1, \dots, v_N) \in V_h^N \\ p &= (p_1, \dots, q_N) \in Q_h^N \\ m &= (u_0, g_1, \dots, g_N) \in M \end{aligned}$$

The sequence of variational problems (12) is reformulated as an operator equation combining all the time steps,

$$(13) \quad \mathcal{F}(m, y) = \sum_{k=0}^{N-1} \{ \mathcal{F}_{k,\Omega}(m, y) + \mathcal{F}_{k,\Gamma_D}(m, y) \} = 0,$$

where $\mathcal{F}_{k,\Omega} : Y \rightarrow Y^*$ is the operator combining all the volume integrals in (12), i.e.

$$\begin{aligned} \mathcal{F}_{k,\Omega}(m, y; v, q) &= \int_{\Omega} \left(\frac{u^{k+1} - u^k}{\delta t} \right) \cdot v^{k+1} + \nu \nabla u^{k+\theta} : \nabla v^{k+1} dx \\ (14) \quad &+ \int_{\Omega} (u^{k+\theta} \cdot \nabla) u^{k+\theta} \cdot v^{k+1} dx \\ &+ \int_{\Omega} q^{k+1} \nabla \cdot u^{k+1} + p^{k+1} \nabla \cdot v^{k+1} dx, \end{aligned}$$

for all $(v, q) = \{(v_k, q_k)\}_{k=1}^N \in Y^*$, for $k = 0, \dots, N-1$. Note that that this part only involves the initial data u_0 from m . The operator $\mathcal{F}_{k,\Gamma_D} : Y \times M \rightarrow Y^*$ combines all the boundary integrals in (12) and reads

$$\begin{aligned} \mathcal{F}_{k,\Gamma_D}(m, y; v, q) &= - \int_{\Gamma_D} \left(\nu \frac{\partial u^{k+\theta}}{\partial n} - p^{k+1} n \right) \cdot v^{k+1} ds \\ (15) \quad &- \int_{\Gamma_D} \left(\theta \nu \frac{\partial v^{k+1}}{\partial n} - q^{k+1} n \right) \cdot (u^{k+1} - g^{k+1}) ds \\ &+ \int_{\Gamma_D} \frac{\nu \sigma}{h} (u^{k+1} - g^{k+1}) \cdot v^{k+1} ds. \end{aligned}$$

The first integral in (15) arises when the integration by parts formula is applied to (11), and the integral would vanish if the Dirichlet boundary condition (2b) were strongly imposed on the space V_h . The remaining terms are added to obtain a variational problem that is consistent and stable, see e.g. [3, 4]. The form (15) is linear and

symmetric, and positive definite provided that the parameter σ is sufficiently large. We must also require $\theta > 0$ to apply the Nitsche method.

The numerical examples in section 4 use two common finite element pairs: P2-P1 (Taylor-Hood) and P1-P1. The lowest order discretisation does not satisfy the LBB conditions, and hence requires stabilisation. We used the stabilisation $-\beta h^2 (\nabla p, \nabla q)_\Omega$ where h is the local mesh element size and $\beta = 10^{-3}$ is the stabilisation coefficient.

3.3.2. Adjoint equations. The adjoint equations are used to efficiently compute the functional derivative $dJ/dm : Y \times M \rightarrow M^*$, at a cost of roughly one linearised Navier-Stokes solve.

To derive the adjoint equations consider the Navier-Stokes equations in the operator form $\mathcal{F}(m; y) = 0 \in Y^*$ and a functional $J(y, m) \in \mathbb{R}$. The total derivative of the functional in direction \tilde{m} is

$$(16) \quad \left\langle \frac{dJ}{dm}, \tilde{m} \right\rangle_{M^*, M} = \left\langle \frac{\partial J}{\partial y}, \frac{dy}{dm} \tilde{m} \right\rangle_{Y^*, Y} + \left\langle \frac{\partial J}{\partial m}, \tilde{m} \right\rangle_{M^*, M}.$$

Evaluating (16) directly is challenging because computing $dy/dm(m) \in \mathcal{L}(M, Y)$ is computationally expensive. The adjoint approach eliminates this term by taking the derivative of the PDE equation

$$(17) \quad \frac{\partial \mathcal{F}}{\partial y} \frac{dy}{dm} + \frac{\partial \mathcal{F}}{\partial m} = 0.$$

and substituting it into (16):

$$(18) \quad \left\langle \frac{dJ}{dm}, \tilde{m} \right\rangle_{M^*, M} = - \left\langle \frac{\partial \mathcal{F}}{\partial m} \tilde{m}, \left(\frac{\partial \mathcal{F}}{\partial y} \right)^{-*} \frac{\partial J}{\partial y} \right\rangle_{Y^*, Y} + \left\langle \frac{\partial J}{\partial m}, \tilde{m} \right\rangle_{M^*, M}.$$

The functional derivative is then computed in two steps:

1. Compute the adjoint solution $\lambda \in Y$ by solving the adjoint PDE

$$(19) \quad \left(\frac{\partial \mathcal{F}}{\partial y} \right)^* \lambda = - \frac{\partial J}{\partial y}$$

2. Evaluate the derivative with

$$(20) \quad \frac{dJ}{dm} = \left(\frac{\partial \mathcal{F}}{\partial m} \right)^* \lambda + \frac{\partial J}{\partial m}$$

The computational expensive part is the solution of (19), which involves the solution of a linear PDE.

The adjoint equations (19) can be derived before or after the discretisation of the Navier-Stokes equations. Here, we chose the discretise-then-adjoint approach, which has the advantage that the discretised derivative is the exact derivative of the discretised system. The alternative approach does not guarantee this, and simple descent methods like algorithm 1 may fail, as demonstrated in [14]. As a consequence a more robust optimisation algorithm would need to be implemented.

The adjoint system (19) for the discretised Navier-Stokes operator (13) is

$$(21) \quad \left\langle \left(\frac{\partial \mathcal{F}}{\partial y} \right)^* \lambda, w \right\rangle = \left\langle \left(\frac{\partial \mathcal{F}}{\partial y} \right) w, \lambda \right\rangle = - \left\langle \frac{\partial J}{\partial u}, v \right\rangle,$$

for all $w = (v, q) \in Y$. Note that the derivative of the regularisation term in the functional vanishes because it does not depend on the state. Since the adjoint operator is linear, it can be written in matrix form:

$$(22) \quad \left(\frac{\partial \mathcal{F}}{\partial y} \right)^* = \begin{pmatrix} \frac{\partial \mathcal{F}_0}{\partial y^1} & 0 & 0 & \cdots \\ \frac{\partial \mathcal{F}_1}{\partial y^1} & \frac{\partial \mathcal{F}_1}{\partial y^2} & 0 & \ddots \\ 0 & \frac{\partial \mathcal{F}_2}{\partial y^2} & \ddots & \ddots \end{pmatrix}^* = \begin{pmatrix} \frac{\partial \mathcal{F}_0^*}{\partial y^1} & \frac{\partial \mathcal{F}_1^*}{\partial y^1} & 0 & \cdots \\ 0 & \frac{\partial \mathcal{F}_1^*}{\partial y^2} & \frac{\partial \mathcal{F}_2^*}{\partial y^2} & \ddots \\ 0 & 0 & \ddots & \ddots \end{pmatrix}$$

or more compactly,

$$\left(\frac{\partial \mathcal{F}}{\partial y^k} \right)^* \lambda = \begin{cases} \left(\frac{\partial \mathcal{F}_{k-1}}{\partial y^k} \right)^* \lambda_k + \left(\frac{\partial \mathcal{F}_k}{\partial y^k} \right)^* \lambda_{k+1} & \text{if } k < N \\ \left(\frac{\partial \mathcal{F}_{k-1}}{\partial y^k} \right)^* \lambda_k & \text{if } k = N. \end{cases}$$

The system (22) is upper-triangular, hence the adjoint (21) is solved by backwards substitution. Written explicitly, the volume integrals in the equation for λ_k , $k = 0, \dots, N$, are

$$(23) \quad \begin{aligned} & \int_{\Omega} \left(\frac{\lambda^k - \lambda^{k+1}}{\delta t} \right) \cdot v^k dx + \int_{\Omega} \nu \nabla \lambda^{k+\tilde{\theta}} : \nabla v^k dx \\ & + \theta \int_{\Omega} (u^{k-1+\theta} \cdot \nabla) v^k \cdot \lambda_u^k dx + \theta \int_{\Omega} (v^k \cdot \nabla) u^{k-1+\theta} \cdot \lambda_u^k dx \\ & + \tilde{\theta} \int_{\Omega} (u^{k+\theta} \cdot \nabla) v^k \cdot \lambda_u^{k+1} dx + \tilde{\theta} \int_{\Omega} (v^k \cdot \nabla) u^{k+\theta} \cdot \lambda_u^{k+1} dx \\ & + \int_{\Omega} \lambda_p^k \nabla \cdot v^k + q^k \nabla \cdot \lambda_u^k dx, \end{aligned}$$

with $(\lambda_u^k, \lambda_p^k) = \lambda^k$, $\tilde{\theta} = 1 - \theta$ and setting $\lambda_u^{N+1} = 0$. Similarly, the boundary integrals are

$$(24) \quad \begin{aligned} & - \int_{\Gamma_D} \left(\nu \frac{\partial v^k}{\partial n} \right) \cdot \lambda_u^{k+\tilde{\theta}} ds + \int_{\Gamma_D} (q^k n) \cdot \lambda_u^k ds \\ & - \int_{\Gamma_D} \left(\theta \nu \frac{\partial \lambda_u^k}{\partial n} - \lambda_p^k n \right) \cdot v^k ds + \int_{\Gamma_D} \frac{\nu \sigma}{h} v^k \cdot \lambda_u^k ds. \end{aligned}$$

The adjoint equations are solved backwards in time, starting from a zero final condition. The timestepping scheme is the same θ -scheme as for the forward discretisation, but with a modified advective velocity. The homogeneous Dirichlet boundary conditions on the controlled surfaces are enforced with a Nitsche like approach.

3.4. Implementation and verification. The Navier-Stokes solver was implemented in the FEniCS finite element framework [20]. The adjoint solver was automatically derived via the algorithmic differentiation tool dolfin-adjoint [9]. The correctness of the adjoint equations, and the resulting derivatives of the goal functional, were verified using the Taylor remainder convergence test. This test checks that for a sufficiently smooth functional \hat{J} , a correct implementation should satisfy

$$(25) \quad \left| \hat{J}(m + h\delta m) - \hat{J}(m) - h \left\langle \frac{d\hat{J}(m)}{dm}, \delta m \right\rangle \right| = O(h^2),$$




FIG. 3. The computational domain for the 2D example. The grey area indicates the extended area used to generate the measurement data.

Parameter	Symbol	Value
Viscosity	ν	3.5
Model timestep	Δt	0.004625 s
End time	T	0.555 s
Time discretisation	θ	0.5
Spatial discretisation		$P1-P1$
Mesh triangles		20,989
Nitsche coefficient	σ	100.0
Number of observations	N	16
Regularisation	$\alpha = \gamma$	10^{-5}

TABLE 1

The settings for the 2D aneurysm flow reconstruction. The first parameters specify the model setup, while the final two parameters configure the data assimilation.

where δm is a control perturbation and $h > 0$ the perturbation size. The Taylor remainder convergence test was performed for different m and random directions δm . Second order convergence was consistently observed, giving confidence that the adjoint implementation is correct.

4. Experiments. In this section, the data assimilation is applied to two different experiments. The first experiment uses an aneurysm-like domain in 2D with known exact solution (section 4.1). The second experiment aims to reconstruct the flow conditions in a real geometry in 3D with observations from an 4D MRI scan (section 4.2).

The implementation and files needed to reproduce the results of this section are available on bitbucket: https://bitbucket.org/biocomp/navier_stokes_data_assimilation. This website contains a Readme file with instructions for the installation and how to reproduce the paper results.

4.1. 2D Aneurysm. This experiment tests the variational data assimilation under idealised conditions where the blood flow to be reconstructed is known a priori. This is used to study the robustness of the reconstruction against incomplete data (both in space and time), noise and the choice of the regularisation. We also compare two different types of observation operators.

The computational domain, shown in figure 3, resembles a blood vessel bifurcation

with an aneurysm in 2D. The observations were generated with the same numerical model that was used in the data assimilation procedure. That is, the Navier-Stokes equations was first solved on an extended domain (including the gray area in figure 3) and the velocity solution used to generate the observations. For this setup, the initial velocity was set to zero. On the inlet and right outlet boundaries a parabolic velocity profiles was enforced with peak values of 1000 mm/s (inlet) and 870 mm/s (right outlet), multiplied by $\sin(\pi(1-t)^3)$ to obtain a pulse like flow pattern¹. The simulation started from a zero velocity at $t = 0$ and was terminated after the peak velocity at $t = 0.629$. The remaining model settings are listed in table 1.

The observation operator \mathcal{T} was applied to the resulting velocity to obtain $N = 16$ observations. We compared two observation operators: the *instantaneous observation operator*, which takes instantaneous measurements at evenly distributed times t_n :

$$(26) \quad \mathcal{T}^{\text{inst}} u := \mathcal{R}_{\Omega_{\text{obs}}} u(t_n),$$

where $\mathcal{R}_{\Omega_{\text{obs}}}$ restricts the velocity to the observation domain (the white area in figure 3). The restriction avoids the “inverse crime” and simulates the incompleteness of real measurement data. The *time-averaging observation operator* takes pointwise time-averaged observations over each observation time interval:

$$(27) \quad \mathcal{T}^{\text{avg}} u := \frac{1}{t_n - t_{n-1}} \int_{t_{n-1}}^{t_n} \mathcal{R}_{\Omega_{\text{obs}}} u(t) dt.$$

The data assimilation was then applied to recover the original flow from the observations. The reconstructions were performed on the restricted domain Ω_{obs} , that is without any knowledge about the geometry of the extended domain. Furthermore, the outflow Dirichlet boundaries were swapped between the data generation and reconstruction to further avoid the “inverse crime”. The optimisation was terminated when the relative change of the functional in one iteration dropped below $|J(x_k) - J(x_{k-1})|/|J(x_0)| \leq 10^{-4}$ or if the number of iteration exceeded 100.

The results of the data assimilation with $\mathcal{T}^{\text{inst}}$ and \mathcal{T}^{avg} are shown in figures 4 and 5 (left column), respectively. The first three plots show the observed and reconstructed velocities at $t = 0.296$ s. Visually, the observed and assimilated velocities agree well.

Since the true velocity is known from the initial simulation, the reconstruction error can also be quantified more rigorously. We define following two error measures: the first measures the relative error of the reconstructed velocity in the aneurysm

$$(28) \quad \mathcal{E}_{\Omega_{\text{ane}}} = \frac{\|u_{\text{true}} - u\|_{\Omega_{\text{ane}} \times (0, T]}}{\|u_{\text{true}}\|_{\Omega_{\text{ane}} \times (0, T]}},$$

where u_{true} is the true velocity. The second measures the relative error of the reconstructed wall shear stress on the aneurysm wall, motivated by the fact that this an important diagnostic value in blood flow simulations:

$$(29) \quad \mathcal{E}_{\text{WSS}} = \frac{\|\text{WSS}(u_{\text{true}}) - \text{WSS}(u)\|_{\Gamma_{\text{ane}} \times (0, T]}}{\|\text{WSS}(u_{\text{true}})\|_{\Gamma_{\text{ane}} \times (0, T]}},$$

with $\text{WSS}(u) = |\sigma n - (\sigma n \cdot n)n|$ and $\sigma = \rho(-pI + \nu(\nabla u + (\nabla u)^T))$ with $\rho = 1060$ kg/m³. The timeplots in figures 4 and 5 (left column) visualise these error measures over the simulation period. The results show a good agreement throughout the simulation period.

¹Note that real flow in cerebral arteries will never go to zero, but rather pulsate between around 0.5 m/s and one third of 0.5 m/s.

4.1.1. Sensitivity of reconstruction with respect to parameter changes.

In this section, we investigate how the quality of the reconstruction depends on noise in the observations and the choice of the reconstruction parameters, such as the amount of regularisation. Since the exact solution is known for this example, we can visualise the reconstructed and the “true” velocities and compute the error measures (28) and (29). The following tests are based on the configuration listed in table 1, and in each test one parameter is varied and the quality of the reconstruction investigated.

Noisy observations (figures 4 and 5). Pointwise Gaussian white noise was added to the observations with zero mean and varying magnitude. This type of noise is not expected in real observations, in particular because it depends on the numerical mesh. Nevertheless, we consider it as a suitable benchmark setup. The results of the reconstruction for different signal to noise ratios are shown in figure 4, and for the instantaneous observation operator $\mathcal{T}^{\text{inst}}$ and figure 5 for the time-averaging observation operator \mathcal{T}^{avg} . With increasing level of noise, the optimised functional value $\mathcal{J} + \mathcal{R}$ increases, because of the increased difference between reconstructed and observed velocity. Nevertheless, the error measures remain small, showing that the reconstruction works reliable even for high noise to signal ratios. Overall, the assimilated flows and metrics agree well for all noise levels, and one can conclude that the reconstruction is little affected by this type of noise.

Regularisation (figures 6 and 7). The regularisation terms (5) enforce “smoothness” on the control functions. Hence the choice of the regularisation coefficients α and γ could have a strong influence on the assimilation results. For the experiments, we varied α and γ coefficients simultaneously to retain the balance between the two regularisation terms. The results for different regularisation values are shown in figures 6 and 7 for $\mathcal{T}^{\text{inst}}$ and \mathcal{T}^{avg} , respectively. The reconstruction works well for values between 10^{-3} and 10^{-5} , but the quality starts to reduce visibly when $\alpha = \gamma > 10^{-2}$. For the case $\alpha = \gamma = 1$, the assimilated velocity is significantly lower than the true velocity, because the strong regularisation enforces spatially and temporally nearly constant controls.

Data sparsity (figures 8 and 9). Another important question is how many observations (N in (3)) are required to accurately reconstruct the blood flow. To address this question, the data assimilation was repeated with varying number of observations N . The base setup (left column in figures 4 and 5) used $N = 16$ and the results for $N = 4, 8$ and 32 are shown in figures 8 and 9 for $\mathcal{T}^{\text{inst}}$ and \mathcal{T}^{avg} , respectively.

With 4 observations the quality of the reconstruction suffers visibly, mostly at the beginning and the end of the simulation times. The time-averaging observation operator yields good results already with 8 observations, while the instantaneous observation operator requires 16 observations to yield an accurate reconstruction. The differences between 16 to 32 observations are minimal for both observation operators.

Choice of controlled outflow boundary (figure 10). In the problem definition (2b), we made a choice to control the outflow on Γ_{out_1} , and to enforce a no-stress condition on Γ_{out_2} . It is therefore natural to check if the reconstruction works well also if Γ_{out_2} is controlled and a non-stress condition is applied on Γ_{out_1} . The results for this setup are shown in figure 10. The reconstruction is similarly good as in the base setup, indicating that the assimilation is not impacted significantly by the choice of the controlled boundaries. Nevertheless, this choice might be more significant for other setups, in particular if one of the outflows is in close proximity to the aneurysm.

Parameter	Symbol	Value
Viscosity	ν	7.5
Model timestep	Δt	0.004625 s
End time	T	0.629 s
Time discretisation	θ	1.0
Spatial discretisation		P1-P1
Dimension of spatial discretisation		184,464
Nitsche coefficient	σ	100.0
Number of observations	N	16
Regularisation parameter	$\alpha = \gamma$	10^{-5}

TABLE 2

The numerical settings for the reconstruction of blood flow from 4D MRA measurements. The first parameters specify the model setup, while the final two parameters configure the data assimilation.

4.2. Flow reconstruction in an aneurysm from 4D MRA measurements.

In this experiment, the variational data assimilation was applied to reconstruct the blood flow conditions in an artificially introduced aneurysm. The measurements were done using 4D PC-MRA, and are described in detail in [15]. The geometry was reconstructed from an image obtained by time-averaging the observations, using VMTK (www.vmtk.org). The observations were then linearly interpolated onto the resulting mesh nodes.

The numerical settings are listed in table 2. The assimilation terminated after 30 optimisation iterations, when the relative change in one optimisation iteration dropped below 0.09%.

For comparison, we additionally performed a high-resolution, low-viscosity flow simulation with “common” choices for the initial and boundary conditions, but without the data assimilation procedure described in this paper. To avoid spurious effects near the boundary for this simulation, the segmented geometry needed to be extended by artificial straight arteries on the in- and outlets. The high-resolution setup had 20 million DOFs with a Taylor-Hood pressure-corrector scheme and a timestep of $5.9 \cdot 10^{-4}$ s. Womersley boundary conditions were used on the inflow and outflows. The inflow flux $Q_{\Gamma_{\text{in}}}$ was interpolated from averaged observations as $(Q_{\Gamma_{\text{in}}} - Q_{\Gamma_{\text{out}_1}} - Q_{\Gamma_{\text{out}_2}})/3$, the outflow flux on Γ_{out_1} was averaged as $Q_{\Gamma_{\text{out}_1}}/(Q_{\Gamma_{\text{out}_1}} + Q_{\Gamma_{\text{out}_2}})$, and a traction free condition was applied on Γ_{out_2} .

The results for the data assimilation approach and the high-resolution solver are shown in figure 11. The figure shows the data assimilation with the instantaneous observation operator - the results for the averaged observation operator look similar. The high-resolution solution has transient to turbulent behaviour in the aneurysm, while the assimilated solution is laminar. Visually, the assimilated solution fits better to the observed velocity, both in the vessel and the aneurysm areas.

5. Conclusion. This paper presented the application of variational data assimilation to reconstruct transient blood flow from observations such as MRI images. This technique is well known in other scientific fields such as ocean science and meteorology, but has thus far not been applied to 3D transient blood flow reconstruction. Mathematically, the data assimilation problem is an optimisation problem constrained by the Navier-Stokes equations. We derived the reduced formulation and described the numerical solution with a focus on retaining the function spaces in the optimisation to obtain mesh-independent iteration numbers in the optimisation step.

The data assimilation was applied to two examples: first, the reconstruction of blood flow in an idealised blood vessel with known solution. This example was used to demonstrate that the proposed method is robust against user parameters and noisy observations. The second example was based on real 4D MRI measurements in a three-dimensional domain, and the result compared to a high-resolution blood flow simulation.

Even though the considered blood flow model and observation operators are simplified, the presented framework extends naturally to more complex setups. Possible extension is to take into account the movement of the vessel wall, non-Newtonian effects or a more realistic observation operator that reimplements an existing measurement device. Furthermore, the reconstruction controls could be extended, for example to also reconstruct the vessel geometry along with the initial and boundary conditions.

The data assimilation procedure introduces an additional computational burden on the flow reconstruction process - for the discussed examples the data assimilation is typically around 50 times more computationally expensive than a single flow simulation. To keep the computational time feasible, the mesh and time resolutions had to be reduced compared to a single flow simulation study. A simple solution would be to first perform a data assimilation on a coarse setup, and then apply the reconstructed initial and boundary conditions on a high-resolution simulation.

Acknowledgments. This research was supported by The Research Council of Norway through a Centres of Excellence grant to the Center for Biomedical Computing at Simula Research Laboratory, project number 179578, and a FRIPRO grant, project number 251237. Computations were performed on the Abel supercomputing cluster at the University of Oslo via NOTUR projects NN9279K and NN9316K. In addition, the authors would like to thank Jingfeng Jiang and Charles Strother for providing the 4D phase-contrast MRA data, and Lorenz John for comments on the regularisation.

REFERENCES

- [1] V. ALESSANDRO, *Data assimilation for incompressible Navier-Stokes: Merging of images, measurements and numerical results in blood flow simulations*, in Proceedings of the V European Conference on Computational Fluid Dynamics, 2010.
- [2] D. AMBROSI, A. QUARTERONI, AND G. ROZZA, *Modeling of physiological flows*, vol. 5, Springer, 2012, doi:10.1007/978-88-470-1935-5.
- [3] R. BECKER, E. BURMAN, AND P. HANSBO, *A Nitsche extended finite element method for incompressible elasticity with discontinuous modulus of elasticity*, Computer Methods in Applied Mechanics and Engineering, 198 (2009), pp. 3352–3360, doi:10.1016/j.cma.2009.06.017.
- [4] E. BURMAN AND P. HANSBO, *A unified stabilized method for Stokes’ and Darcy’s equations*, Journal of Computational and Applied Mathematics, 198 (2007), pp. 35–51, doi:http://dx.doi.org/10.1016/j.cam.2005.11.022.
- [5] J. R. CEBRAL, F. MUT, J. WEIR, AND C. PUTMAN, *Quantitative characterization of the hemodynamic environment in ruptured and unruptured brain aneurysms*, American Journal of Neuroradiology, 32 (2011), pp. 145–151, doi:10.3174/ajnr.A2419.
- [6] M. D’ELIA, M. PEREGO, AND A. VENEZIANI, *A variational data assimilation procedure for the incompressible Navier-Stokes equations in hemodynamics*, Journal of Scientific Computing, 52 (2012), pp. 340–359, doi:10.1007/s10915-011-9547-6.
- [7] J. M. DOLAN, J. KOLEGA, AND H. MENG, *High wall shear stress and spatial gradients in vascular pathology: a review*, Annals of biomedical engineering, 41 (2013), pp. 1411–1427, doi:10.1007/s10439-012-0695-0.
- [8] Ø. EVJU, K. VALEN-SENDSTAD, AND K.-A. MARDAL, *A study of wall shear stress in 12 aneurysms with respect to different viscosity models and flow conditions*, Journal of biomechanics, 46 (2013), pp. 2802–2808, doi:10.1016/j.jbiomech.2013.09.004.

- [9] P. E. FARRELL, D. A. HAM, S. W. FUNKE, AND M. E. ROGNES, *Automated derivation of the adjoint of high-level transient finite element programs*, SIAM Journal on Scientific Computing, 35 (2013), pp. C369–C393, doi:10.1137/120873558.
- [10] A. V. FURSIKOV, M. D. GUNZBURGER, AND L. S. HOU, *Boundary value problems and optimal boundary control for the Navier–Stokes system: the two-dimensional case*, SIAM Journal on Control and Optimization, 36 (1998), pp. 852–894, doi:10.1137/S0363012994273374.
- [11] G. GOPALAKRISHNAN, B. D. CORNUELLE, I. HOTEIT, D. L. RUDNICK, AND W. B. OWENS, *State estimates and forecasts of the loop current in the Gulf of Mexico using the MITgcm and its adjoint*, Journal of Geophysical Research: Oceans, 118 (2013), pp. 3292–3314, doi:10.1002/jgrc.20239.
- [12] T. GUERRA, A. SEQUEIRA, AND J. TIAGO, *Existence of optimal boundary control for the Navier–Stokes equations with mixed boundary conditions*, Portugaliae Mathematica, 72 (2015), pp. 267–283, doi:10.4171/pm/1968.
- [13] T. GUERRA, J. TIAGO, AND A. SEQUEIRA, *Optimal control in blood flow simulations*, International Journal of Non-Linear Mechanics, 64 (2014), pp. 57–69, doi:10.1016/j.ijnonlinmec.2014.04.005.
- [14] M. GUNZBURGER, *Perspectives in Flow Control and Optimization*, Society for Industrial and Applied Mathematics, 2002, doi:10.1137/1.9780898718720.
- [15] J. JIANG, K. JOHNSON, K. VALEN-SENDSTAD, K.-A. MARDAL, O. WIEBEN, AND C. STROTHER, *Flow characteristics in a canine aneurysm model: a comparison of 4D accelerated phase-contrast MR measurements and computational fluid dynamics simulations*, Medical physics, 38 (2011), pp. 6300–6312, doi:10.1118/1.3652917.
- [16] L. J. JOHN, *Optimal Boundary Control in Energy Spaces*, vol. 24, Verlag der Technischen Universität Graz, 2014, doi:10.3217/978-3-85125-373-3.
- [17] C. T. KELLEY AND E. W. SACHS, *Quasi-Newton methods and unconstrained optimal control problems*, SIAM Journal on Control and Optimization, 25 (1987), pp. 1503–1516, doi:10.1137/0325083.
- [18] Z. KULCSÁR, A. UGRON, M. MAROSFÓI, Z. BERENTEI, G. PAÁL, AND I. SZIKORA, *Hemodynamics of cerebral aneurysm initiation: the role of wall shear stress and spatial wall shear stress gradient*, American Journal of Neuroradiology, 32 (2011), pp. 587–594, doi:10.3174/ajnr.A2339.
- [19] H. LEE, *Optimal control for quasi-Newtonian flows with defective boundary conditions*, Computer Methods in Applied Mechanics and Engineering, 200 (2011), pp. 2498–2506, doi:10.1016/j.cma.2011.04.019.
- [20] A. LOGG, K. A. MARDAL, AND G. N. WELLS, *Automated Solution of Differential Equations by the Finite Element Method*, Springer, 2012, doi:doi:10.1007/978-3-642-23099-8.
- [21] J. NITSCHKE, *Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind*, in Abhandlungen aus dem mathematischen Seminar der Universität Hamburg, vol. 36, Springer, 1971, pp. 9–15, doi:10.1007/BF02995904.
- [22] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer Science & Business Media, 2006, doi:10.1007/978-0-387-40065-5.
- [23] S. RAMALHO, A. MOURA, A. M. GAMBARUTO, AND A. SEQUEIRA, *Sensitivity to outflow boundary conditions and level of geometry description for a cerebral aneurysm*, International Journal for Numerical Methods in Biomedical Engineering, 28 (2012), pp. 697–713, doi:10.1002/cnm.2461.
- [24] S. RAMALHO, A. MOURA, A. M. GAMBARUTO, AND A. SEQUEIRA, *Influence of blood rheology and outflow boundary conditions in numerical simulation of cerebral aneurysms*, Springer, 2013, pp. 149–175, doi:10.1007/978-1-4614-4178-6_6.
- [25] H. SAMADY, P. ESHTEHARDI, M. C. MCDANIEL, J. SUO, S. S. DHAWAN, C. MAYNARD, L. H. TIMMINS, A. A. QUYYUMI, AND D. P. GIDDENS, *Coronary artery wall shear stress is associated with progression and transformation of atherosclerotic plaque and arterial remodeling in patients with coronary artery disease*, Circulation, 124 (2011), pp. 779–788, doi:10.1161/CIRCULATIONAHA.111.021824.
- [26] T. SCHWEDES, S. W. FUNKE, AND D. A. HAM, *An iteration count estimate for a mesh-dependent steepest descent method based on finite elements and Riesz inner product representation*, (2016), arXiv:1606.08069.
- [27] D. A. STEINMAN, *Assumptions in modelling of large artery hemodynamics*, in Modeling of Physiological Flows, Springer, 2012, pp. 1–18, doi:10.1007/978-88-470-1935-5_1.
- [28] H. TAKAO, Y. MURAYAMA, S. OTSUKA, Y. QIAN, A. MOHAMED, S. MASUDA, M. YAMAMOTO, AND T. ABE, *Hemodynamic differences between unruptured and ruptured intracranial aneurysms during observation*, Stroke, 43 (2012), pp. 1436–1439,

- doi:10.1161/STROKEAHA.111.640995.
- [29] J. TIAGO, A. GAMBARUTO, AND A. SEQUEIRA, *Patient-specific blood flow simulations: Setting Dirichlet boundary conditions for minimal error with respect to measured data*, Mathematical Modelling of Natural Phenomena, 9 (2014), pp. 98–116, doi:10.1051/mmnp/20149608.
 - [30] P. TRIVERRO, L. DEDE, A. SEQUEIRA, S. DEPARIS, A. ROBERTSO, AND A. QUARTERONI, *Fluid-structure interaction simulations of cerebral arteries by isotropic and anisotropic constitutive laws*, Computational Mechanics, (2015), doi:10.1007/s00466-014-1117-y.
 - [31] J. XIANG, S. K. NATARAJAN, M. TREMMEL, D. MA, J. MOCCO, L. N. HOPKINS, A. H. SIDDIQUI, E. I. LEVY, AND H. MENG, *Hemodynamic-morphologic discriminants for intracranial aneurysm rupture.*, Stroke, 42 (2011), pp. 144–52, doi:10.1161/STROKEAHA.110.592923.

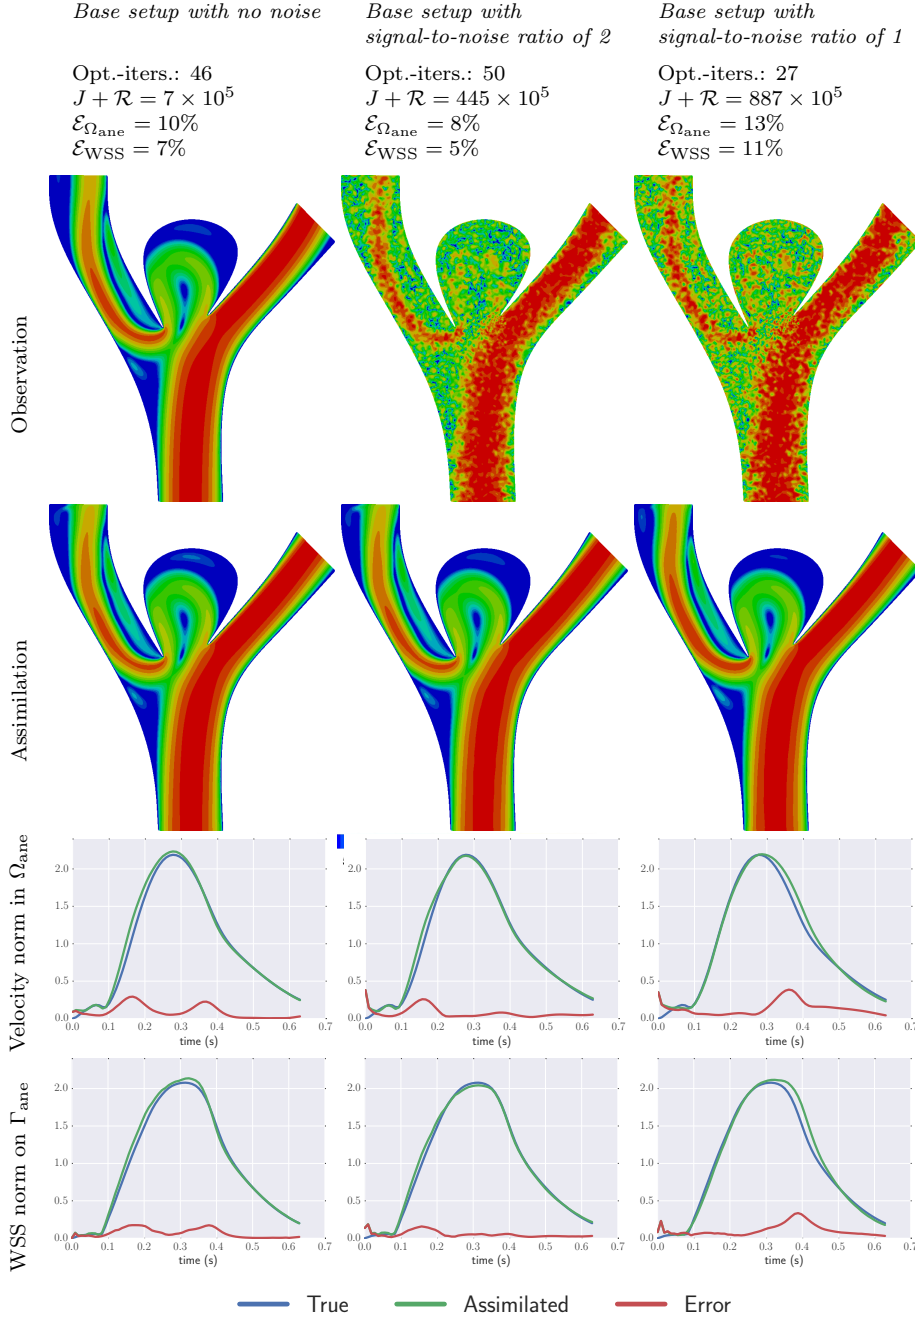


FIG. 4. Results using the *instantaneous observation operator* with pointwise additive Gaussian white noise. The signal-to-noise-ratio was computed as $\|\mathcal{T}^{inst}u_{true}\|^2 / \|\mathcal{T}^{inst}u_{true} - d\|^2$, where d is the noisy data. The snapshots on the top three rows are taken at $t = 0.296s$.

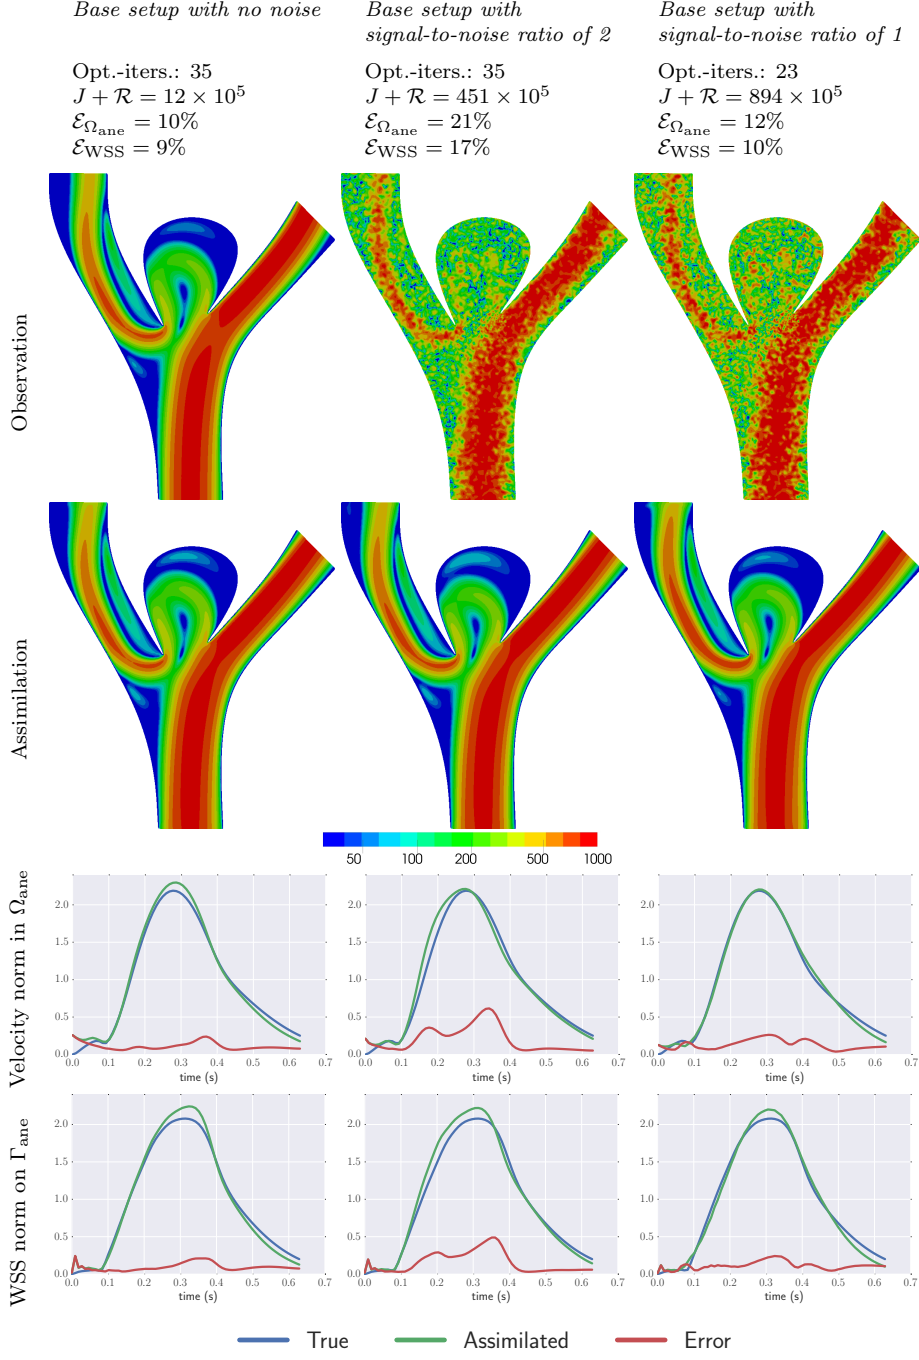


FIG. 5. Results using the **time-averaging observation operator** with pointwise additive Gaussian white noise. The signal-to-noise ratio was computed as $\|\mathcal{T}^{avg}u_{true}\|^2/\|\mathcal{T}^{avg}u_{true} - d\|^2$, where d is the noisy data. The snapshots on the top three rows are taken at $t = 0.296s$.

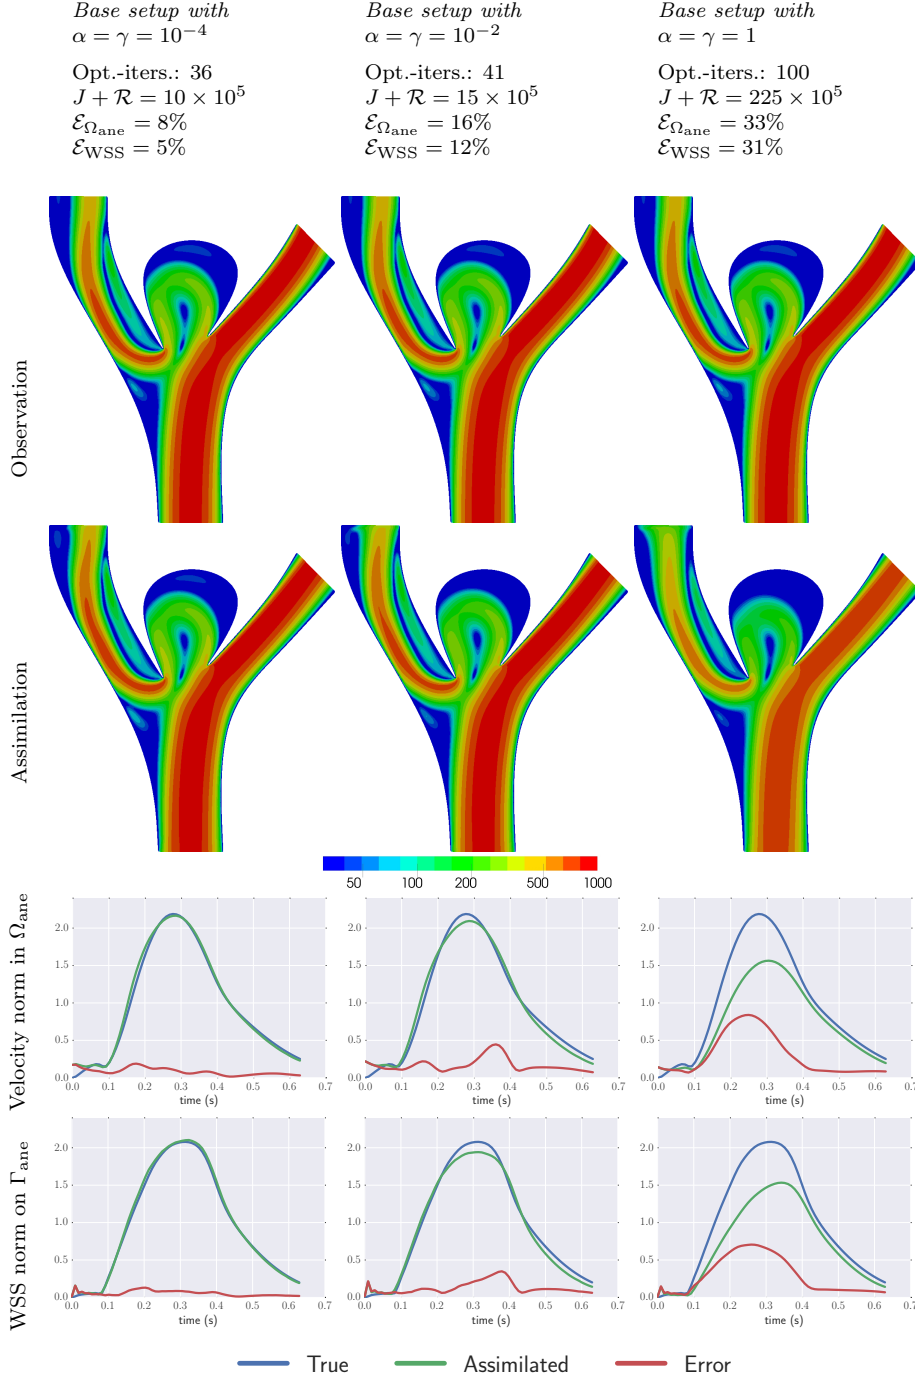


FIG. 6. Results using the *instantaneous observation operator* with varying α and γ regularisation coefficients. The base setup (figure 4, left column) uses $\alpha = \gamma = 10^{-5}$. The snapshots on the top two rows are taken at $t = 0.296\text{s}$.

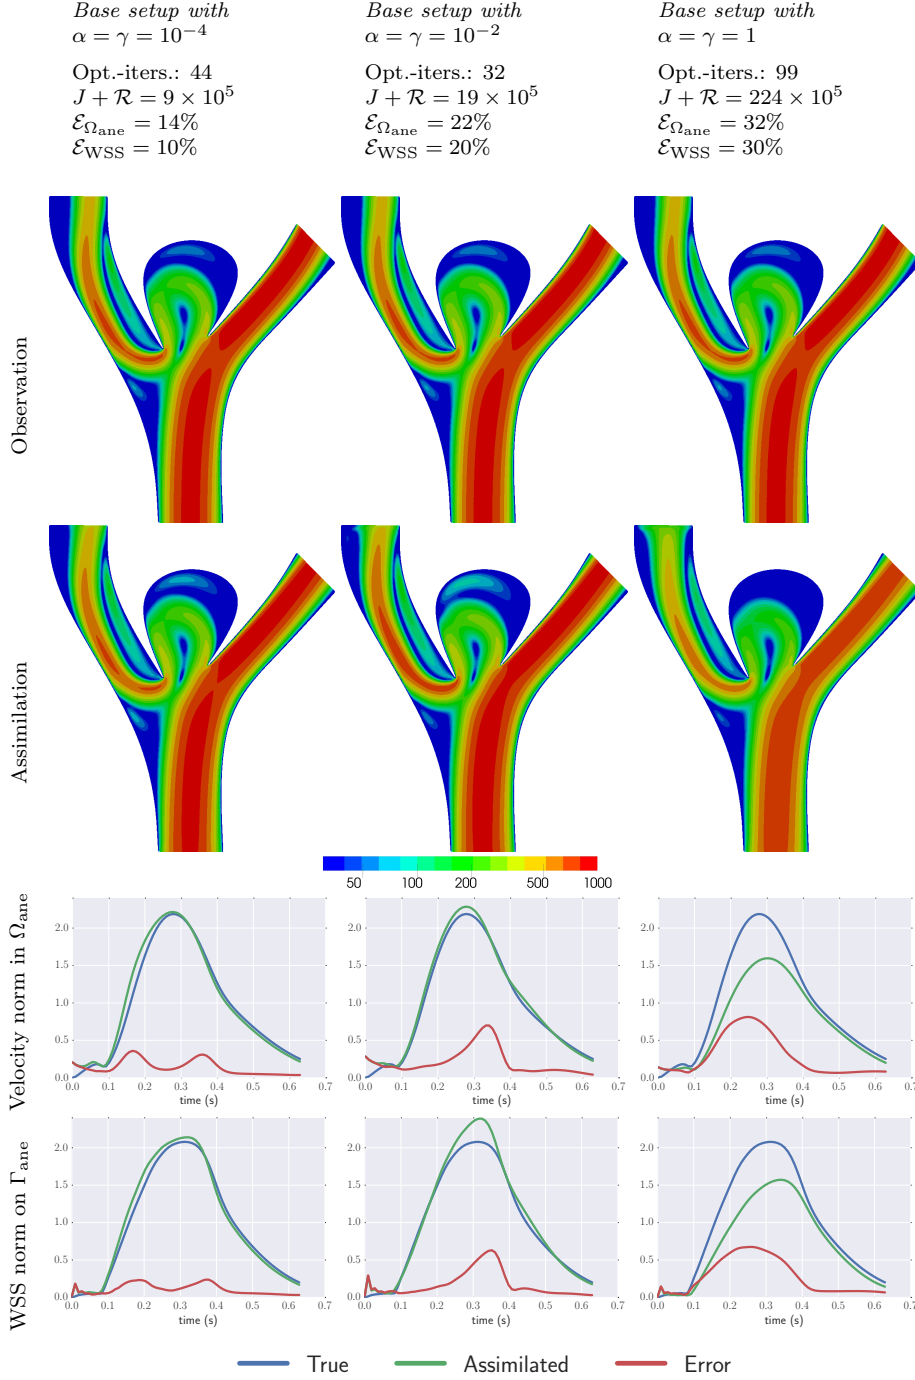


FIG. 7. Results using the **time-averaging observation operator** with varying α regularisation parameters. The base setup (figure 5, left column) uses $\alpha = \gamma = 10^{-5}$. The snapshots on the top two rows are taken at $t = 0.296\text{s}$.

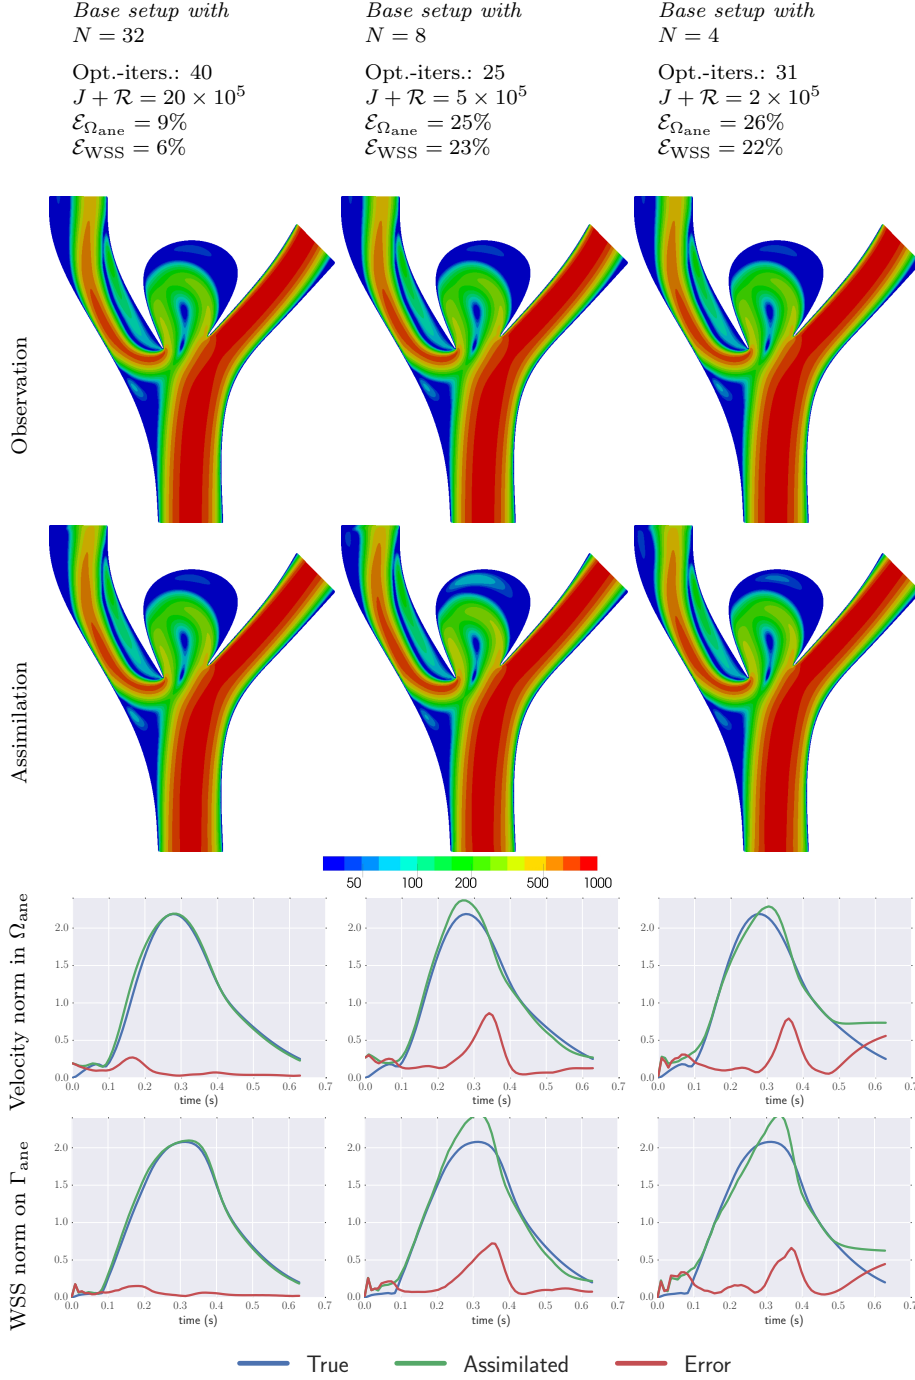


FIG. 8. Results using the *instantaneous observation operator* with varying number of observations. The base setup (figure 4, left column) uses $N = 16$ observations. The snapshots on the top two rows are taken at $t = 0.296s$.

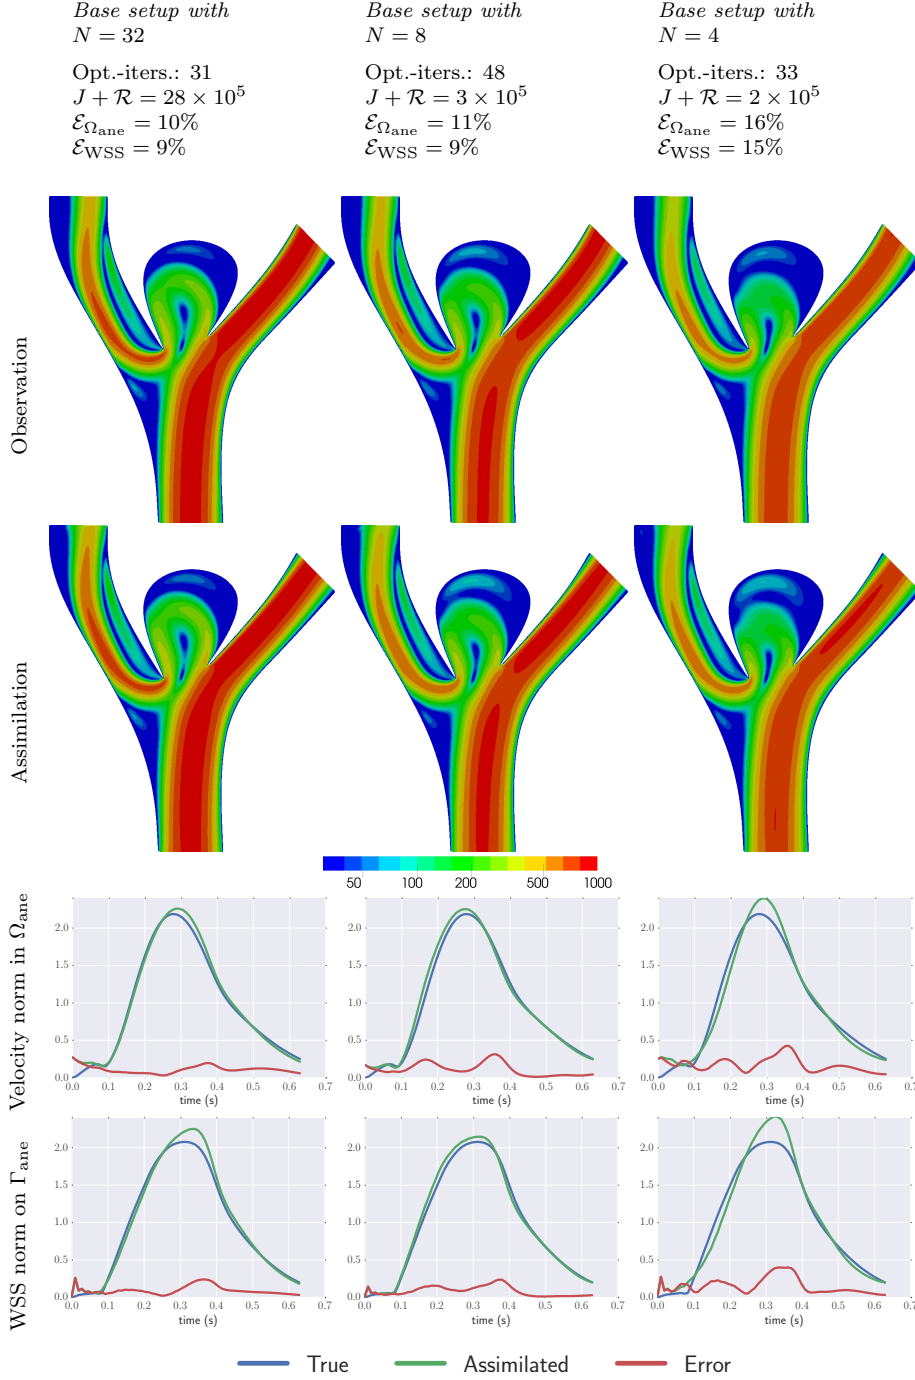


FIG. 9. Results using the **time-averaging observation operator** with varying number of observations. The base setup (figure 5, left column) uses $N = 16$ observations. The snapshots on the top two rows are taken at $t = 0.296\text{s}$.

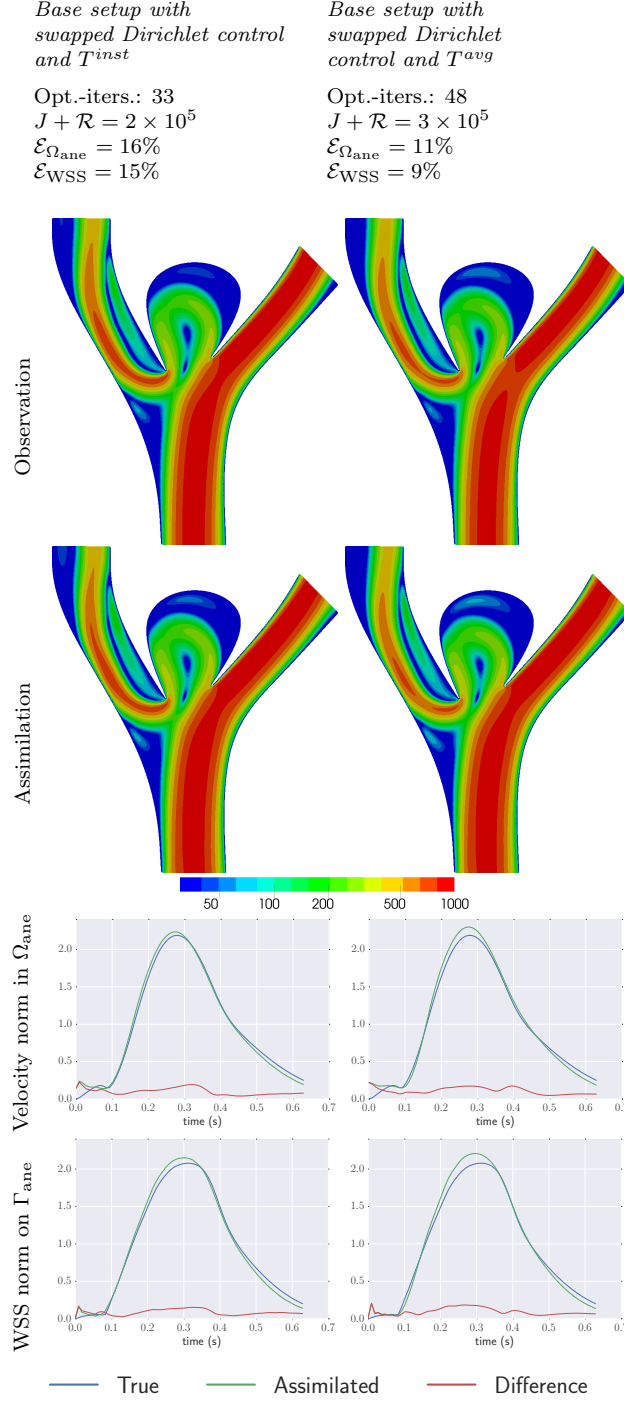


FIG. 10. Assimilation results where the controlled outlets are swapped from the base cases (figures 4 and 5, left columns). The snapshots on the top two rows are taken at $t = 0.296s$.




FIG. 11. Flow speed visualised through a slice of the 3D dog vessel. The snapshots are taken at time $t = 0.296s$.